United States Department of Agriculture

Office of the Chief Information Officer
Network Engineering Division

Performance and Results Management System Network Analysis Project

Phase II Report

June 10, 1999

Executive Summary

An initial evaluation of the WWW-based Performance and Results Management System (PRMS) by the Network Engineering Division (NED) revealed that the principal cause of unacceptable Web sessions is due to a large number of displaced data packets and resulting retransmissions that increase throttling time. While packets were dropped at other locations, the greatest number of retransmissions occurred on the network segment between the Client and the DDR sniffers. Analysis also revealed that the USDA WAN segment of the network, the PRMS client workstation and server components, and the PRMS application did not contribute to unacceptable response times. This report addresses the cause, location, and direction of the PRMS performance problems along with other related issues, such as routing equipment, transmission packet size, and network traffic load factors.

Data analysis in this report shows that Cyclic Redundancy Checksum (CRC) errors on the link between the client and DDR modems are the primary cause of the data packet loss problem, and that the loss occurs in both directions across the DDR link. Hard-to-fix flaws in the public switched telephone network (PSTN) generate the CRC errors. While synchronous technologies such as the PSTN are generally more efficient than asynchronous technologies, this is not the case in the PRMS environment because of the PSTN imperfections. Expanding the DDR rotor will solve port contention problems but will not solve the dropped packet problem and associated unacceptable response times.

Solutions to the packet loss problem include improving line quality, adjusting the Maximum Transmission Unit (MTU), converting to asynchronous technology, and converting to Frame Relay technology. Associated monthly costs should also be considered as part of a thorough cost/benefit analysis. Any chosen solution should be thoroughly tested prior to full-scale implementation.

1.0 Introduction

The Performance and Results Management System (PRMS) is a WWW-based application used to collect performance data from National Resource Conservation Services (NRCS) county offices. The Network Engineering Division (NED) evaluated the network performance of this mission critical application in a project documented as the PRMS *Network Analysis Project—Final Report, May 10, 1999*. On May 11, 1999 a new project began to evaluate further the results of the Phase I project.

Phase I of the PRMS Project resulted in the following conclusions about the Application network performance:

- The principal cause of the unacceptable Web sessions occurs when a large number of packets are displaced during the session. The resulting large number of retransmissions significantly increases the throttling time before the sender releases the next packet, which then incurs unacceptable response times (exceeding 30 seconds).
- Network delay on the LWV DDR portion of the network accounted for the greatest amount of response time during testing.
- Dropped packets occurred during acceptable Web sessions, but the low frequency did not incur unacceptable response times.
- The limited data available did not reveal the exact source of the dropped packets. Possible causes such as network delays and data corruption are explored in this report.
- The greatest number of retransmissions for single packets occurred on the network segment between the Client and the DDR sniffers, although packets were also dropped at other locations.
- The USDA WAN segment of the network is <u>not</u> a major contributing factor to the unacceptable response times for this type of application.
- The PRMS client workstation and server components of the architecture are <u>not</u> contributing factors to the unacceptable response times.
- The operation of the PRMS application is <u>not</u> a contributing factor to the unacceptable response times. The application uses off-the-shelf commercial products to provide the client server connectivity, and no abnormal client-server behavior was observed.

Phase II of the PRMS Project was specifically designed to answer the following questions about the network performance of the PRMS application:

- What is the cause and location of the PRMS performance problem?
- Is the PRMS performance problem a bi-directional transmission issue?
- Does the Dial-on-Demand Routing (DDR) equipment significantly affect PRMS performance?
- Does the application transmission packet size significantly affect PRMS performance?
- Is application performance affected by a network traffic load factor?

2.0 Methodology

2.1 Project Plan

Table 1 describes the tasks developed for the project plan extension. A Gantt chart in Figure 1 presents the schedule for the Phase II project.[1] As the chart shows, the project exceeded the scheduled due date. Two factors contributed to the missed deadline. The first was the time needed to "Re-install DDR Sniffer" and "Extend Test Lab to NED." The second factor was a misconfigured test modem in the NED lab that was discovered on May 27. After correcting the configuration, data collection continued. However, data collection extended beyond May 27 to allow use of a properly configured modem, pushing the final report from June 4 to June 9. (Note: Data collected prior to May 27 was not used to quantify the CRC-error-to-dropped-packet ratio.)

Table 1 PRMS Project Phase II Tasks

Task	Description	Deliverable	Duration	Participation
Re-install DDR Sniffer	Install the DDR sniffer in Kansas City, MO that was removed to assess the performance impact of the device.	Sniffer installed and functioning	4 Days	NED, DDR Operational Staff
Develop New Test Suite & Schedule	Develop a test suite and schedule to better define the nature of the performance problems encountered during PRMS testing	Informal documentation of test procedures and schedule	4 Days	NED, NRCS
Extend Test Lab to NED	Install test lab in the NED lab area. This is to better facilitate the coordination of testing and data collection	Test lab installed and functioning	4 Days	NED, NRCS
Sample Data	Collect statistical data from sniffers, routers, baseline activities and other sources during testing.	Database of sampled application statistics	8 Days	NED, DDR Operational Staff
Analyze Data	Perform statistical analysis of the results of all of the data collection activities	Informal notes, charts, databases and other graphics from the analysis	10 Days	NED
Present Findings	Present the results of the analysis to the PRMS project staff and others as defined by the PRMS project leader	Informal presentation and discussion	1 Day	NED, NRCS
Develop Report	Document the results of the analysis in a formal report.	Formal Report	4 Days	NED

[1] For reference, a Gantt chart describing both Phase I and Phase II of the Project is presented in Appendix A.

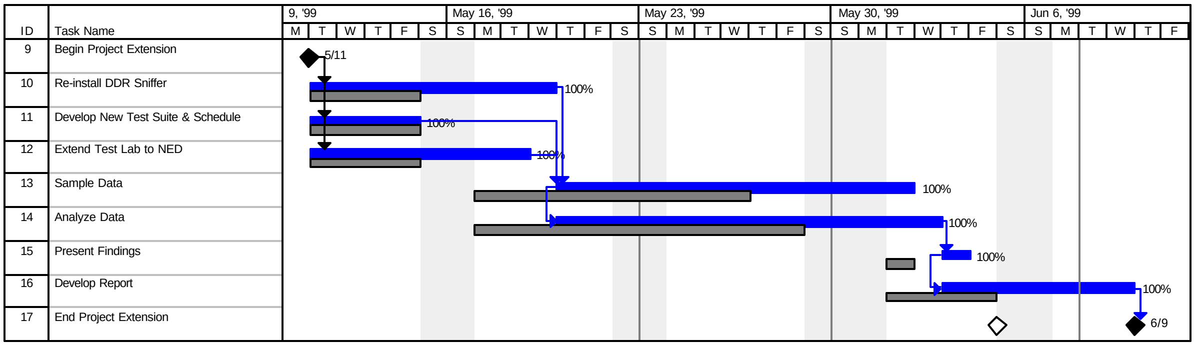

Figure 1 Schedule for PRMS Project Phase II

2.2 Test Environment

The original test environment was modified (see Figure 2) to define further the location and cause of the packet loss problem identified in Phase I of the project. The first modification established a simulated field office environment at the NED office in Fort Collins, Colo., to provide better access to the test environment. As it was in the original office, the simulated client workstation and the network analyzer were the only devices on the LAN. The second modification used a non-PRMS server located in Kansas City, Mo., to eliminate the USDA WAN portion of the original test network and the PRMS application and server, because neither contributed to the performance problem. The non-PRMS server is labeled "ds.usda.gov" in Figure 2. Network analyzers (sniffers) were placed on the LAN adjacent to the PRMS client and on the LAN segment connecting the LWV dial network to the USDA WAN.

In addition to these physical changes, other data collection capabilities were added to the test network. The improved access to the client test environment allowed detailed router information to be recorded from the client router. SNMP level statistics were also collected from the three DDR routers providing a level of detail not available in Phase I testing.

Phase II also used a different test suite. Network pings were used to generate traffic instead of testing downloads of a single 29,304 byte Web page from the PRMS server to the client. (A few Web page transfers were tested to verify the results from Phase I). Using network pings allowed a more controlled data stream across the network and controlled the frequency and size of the packets. All ping testing was done between the PRMS client and the "ds.usda.gov" server. Because PING uses the ICMP protocol, using "ds.usda.gov" as the destination host eliminated the low priority response that network router place on ICMP requests. The testing was coordinated so that the network analyzers captured the client-server traffic during the test, which was performed during normal business hours while the network and other application processes were in operation.

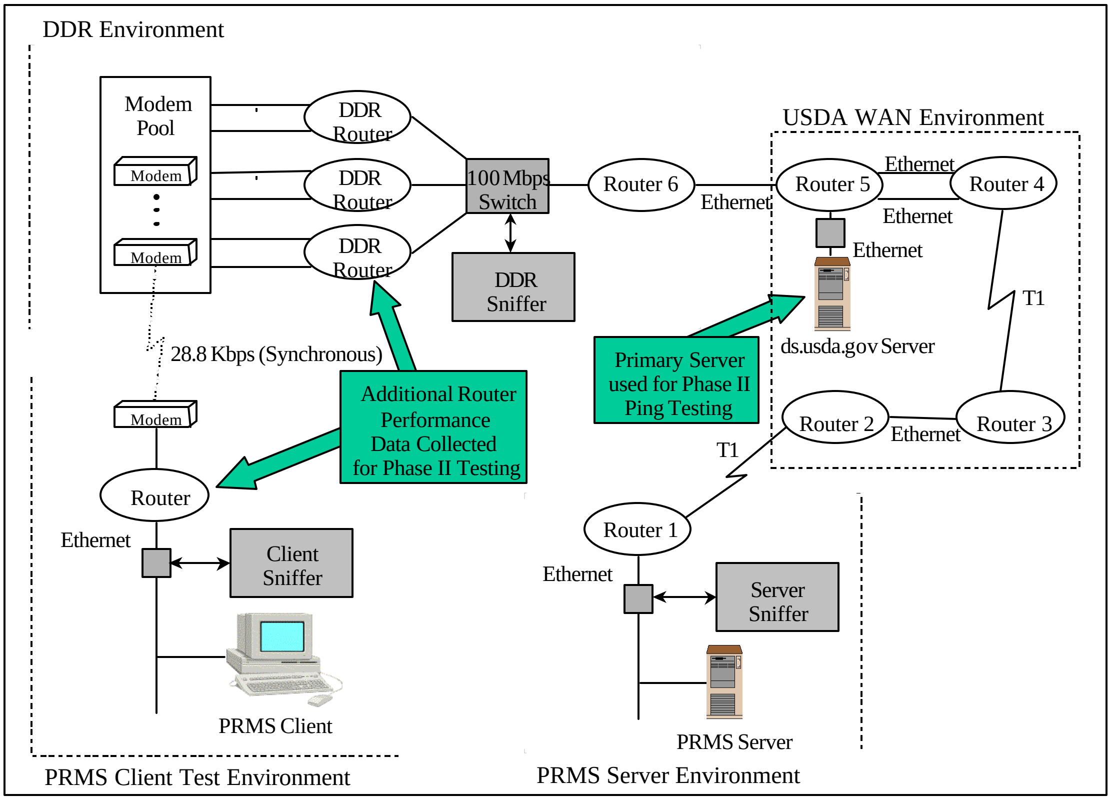

Figure 2 PRMA Phase II Test Network Diagram

Once the test data were collected, they were analyzed using software developed as part of the Resource Planning Methodology (RPM). The software tools processed the traffic data collected by the network analyzers and analyzed the client-server interaction over the network. PRMS developers provided other information about the frequency of failure and server configuration.

3.0 Results

The goal of the Phase II analysis was to answer the questions posed in the introduction of this document. This section addresses each of these questions separately.

3.1 Cause and Location of Problem

Analysis in Phase I identified the location of the problem to be somewhere between the DDR and client environment. To identify the location more specifically, Phase II testing was initiated that showed a high occurrence of Cyclic Redundancy Checksum (CRC) errors on the link between the client and DDR modems. CRC errors occur when the checksum of the data received at the far-end router does not equate. The location of the errors is identified in Figure 3. A high occurrence of packet drops was also recorded along the same link. Figure 4 illustrates a strong positive correlation (89 percent) between displaced packets and CRC errors that occurred on the link between modems.

Noise, gain hits, and other transmission problems on telecommunications links cause CRC errors. These types of problems are typical of remote access solutions that use the public switched telephone network (PSTN). The main problem is the number and quality of switches and PBXs involved in the telephone call. The load on these devices also has a direct impact on a connection's clarity and the quality of the modem signal.

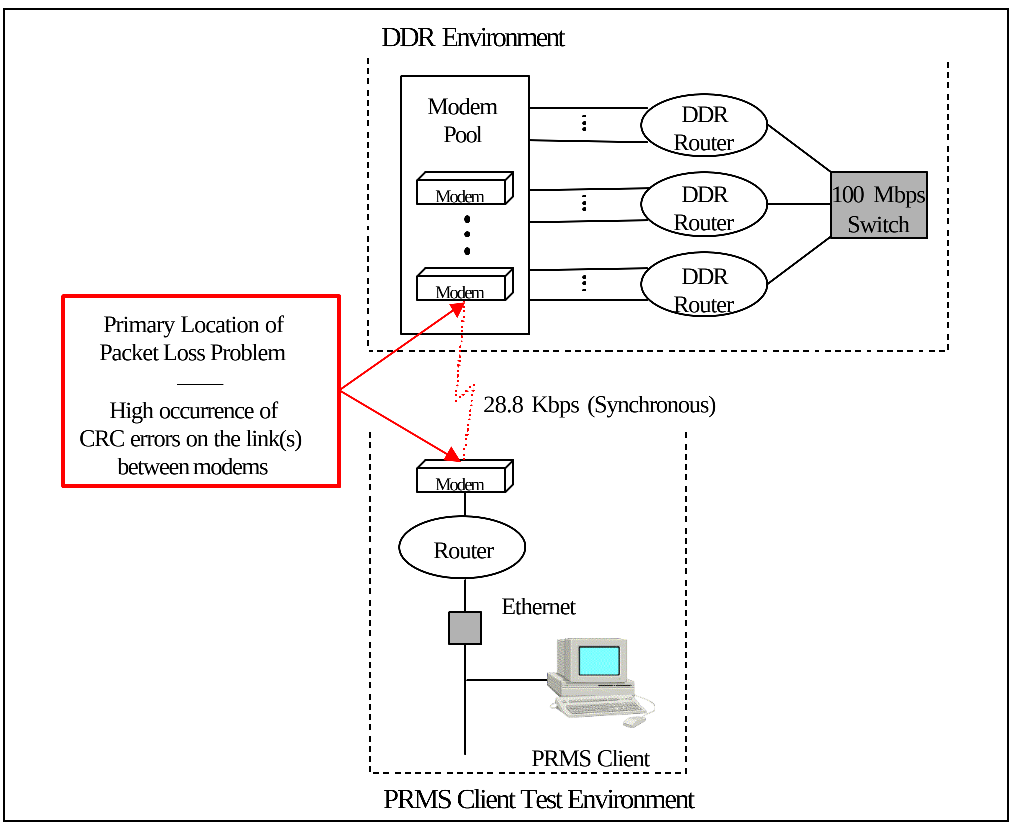

Figure 3 Location of PRMS Packet Loss Problem

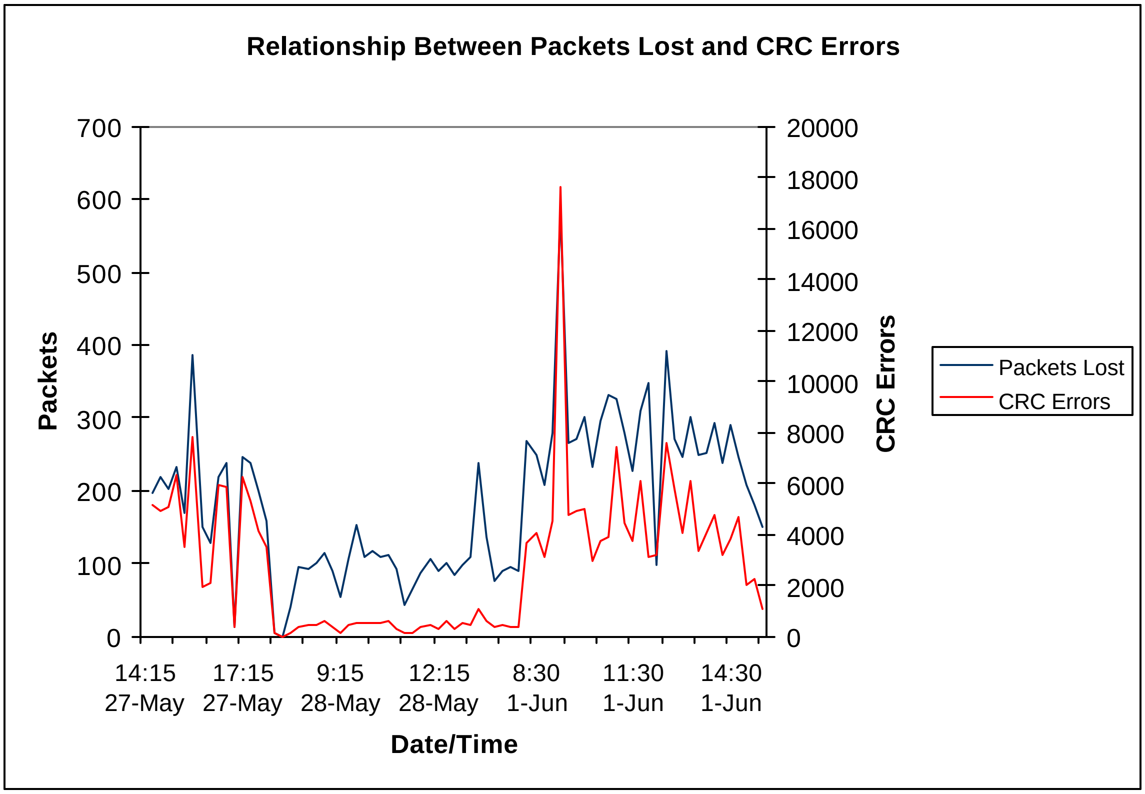

Figure 4 Relationship between Packets Lost and CRC Errors

3.2 Bi-directionality of Problem

In the Phase I analysis, packet loss was only observed in one direction--from the server side of the DDR to the client side. The data transfer traffic flow is predominantly from the PRMS WWW server to client. To test the theory that the direction of the traffic flow causes packet loss, ping testing was run in which the same number and size of packets were transferred in both directions.

Data collected between the link on the client side of the DDR and the client showed a high occurrence of CRC errors. Although data were not available for the DDR routers, the DDR administrator reported they were recording thousands of CRC errors.

3.3 DDR Router Connection

The "Next Steps" section of the Phase I report recommends investigating the possibility of the error only occurring on one of the three DDR routers. The finding that the packet loss is occurring on the links between modems indicates a bad line(s) or modem(s) in the rotor exists. However, the limited testing in Phase II does not provide sufficient data to draw a conclusion. A cursory examination of the limited data indicates the problem is probably not located on a specific DDR router.

The data also show a higher ratio of discards to forwarded datagrams from the third DDR router. This ratio indicates packets are being routed to the third router that it cannot forward. This may be explained by the router's role as the primary "dial-out" device. While outbound connections are being established, packets in the queue may be discarded until the connection is complete. This possibility is presented for informational purposes only. Further study is required to draw a definitive conclusion.

3.4 Data Packet Size

In phase I of the project, larger packet sizes seemed to be dropped more frequently than smaller ones. Ranges of packet sizes were tested across the network to investigate this theory. Test results show a strong correlation between packets dropped and packet size; the larger the packet, the more likely it was to be dropped. Table 2, using data shown in the Appendix B, shows how the success rate of the ping testing decreased as the packet size increased.

For reference, the success rate for all packet sizes on a similar asynchronous connection exceeds 80 percent. The success rate for all packet sizes on a multi-hop Frame Relay network typically exceeds 99 percent.

Table 2 Ping Testing Results Summary

Packet Size (bytes)	Success Rate (%)	Round Trip Time	
		Average (msec)	Median (msec)
1	89.8	485	195
500	66.6	750	723
1000	56.9	865	772
1400	55.1	1025	1015

3.5 Load Factors

Original reports from users of the PRMS application indicated that the rate of dropped packets increased during "busiest" times of the day. The increased drop rate also occurred during phase II testing. In general, the packet loss problem occurred more frequently during peak usage (load) times. The worst time was in the afternoon, improving overnight and in the early morning. However, errors occurred on the modem link at all times of the day.

Figure 4 shows these error rates. May 28 shows a decrease in the number of occurrences of the packet loss problem. May 28 was a Friday before a three-day weekend, which historically is a low usage day. Compared to Tuesday, June 1, a normal business day, one sees a marked increase in the frequency of the packet loss problem.

4.0 Conclusions

The results of Phase II of the PRMS Network Application Performance Analysis Project show that CRC errors on the link between the client and DDR modems are the primary cause of the data packet loss problem, and that the loss occurs in both directions across the client to DDR link. Imperfections in the public switched telephone network (PSTN) generate the CRC errors, and no simple solution exists to correct the PSTN problems. While synchronous technologies, such as the PSTN, are generally more efficient than asynchronous technologies, this is not true in the PRMS environment because of the PSTN flaws. Therefore, solutions using synchronous technologies should be avoided. Expanding the DDR rotor will solve the port contention problem, but the data indicate this solution is unlikely to solve the dropped packet problem and the associated unacceptable response times.

5.0 Discussion and Recommendations

5.1 Possible Solutions

Phase I and II testing and analysis present several possible solutions. Table 2 shows these solutions as well as some ranking criteria. The values assigned to the criteria are assigned based on the information presented in the Phase I and II reports and the technical experience of NED staff rather than statistical information. This information is presented to assist the PRMS staff in evaluating the solutions available to solve the original performance problem. (Note that the cost column of the table represents the estimated magnitude of the cost of implementation and does not include monthly recurring costs.)

Table 3 Possible Solutions to the Packet Loss Problem

Alternative	Description	Cost	Level of Effort	Probability of Solving Port Contention Problem	Probability of Solving Dropped Packet Problem
Improve Line Quality	Work with service provider to improve the quality of the lines connected to the DDR rotor	LOW	LOW	LOW	LOW
Adjust Maximum Transmission Unit (MTU)	Decrease the MTU size on the DDR network link. This will reduce packet sizes and potentially improve the transmission success rate.	LOW	MED	LOW	LOW to MED
Expand DDR rotor	Increase the size of the DDR rotor by adding lines, modems and routers	HIGH	MED	HIGH	LOW
Convert to Asynchronous technology	Modify the modems to operate in Asynchronous mode. Requires new routers to support the interface requirement.	HIGH	HIGH	LOW	MED
Convert to Frame Relay technology	Use Frame relay instead of the dial solution. This conversion is currently underway.	MED to HIGH	HIGH	HIGH	HIGH

5.2 Recommendations

A more thorough cost/benefit analysis that includes the associated monthly costs, as well as the factors presented in Table 4, should be done before selecting a solution. The chosen solution should also be thoroughly tested before full-scale implementation. Upon implementation, the conversion to Frame Relay technology should be expedited to solve both port contention as well as performance problems. When the DDR is in use, examine modem configurations and functionality as well as DDR router performance to ensure they do not impede network performance.

Appendix A

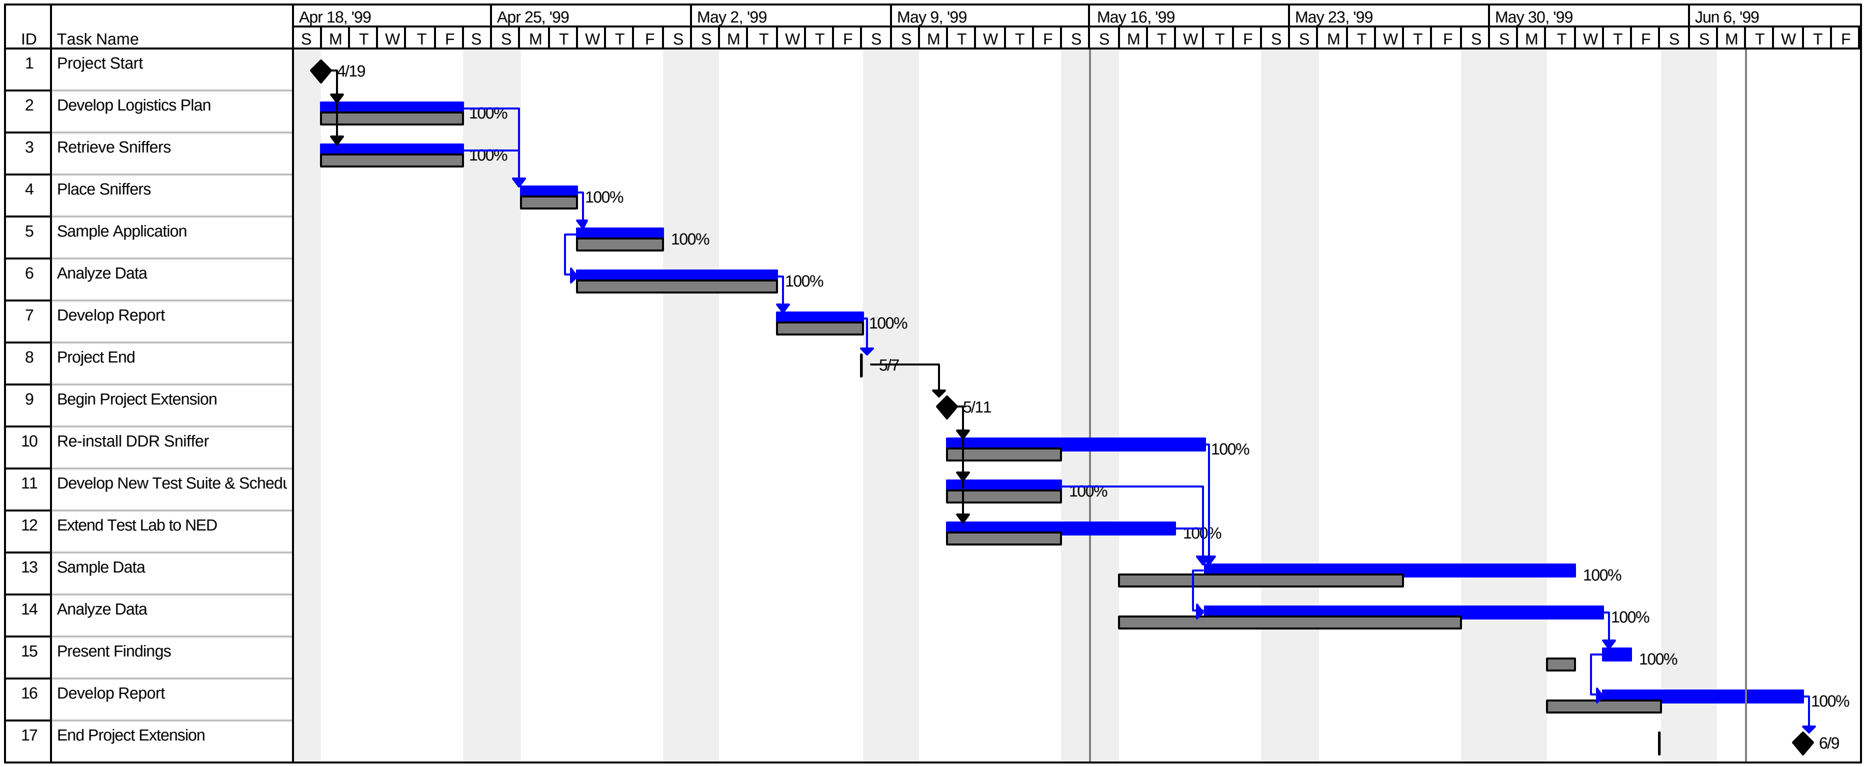

Figure 5 Complete PRMS Project Gantt Chart

Appendix B

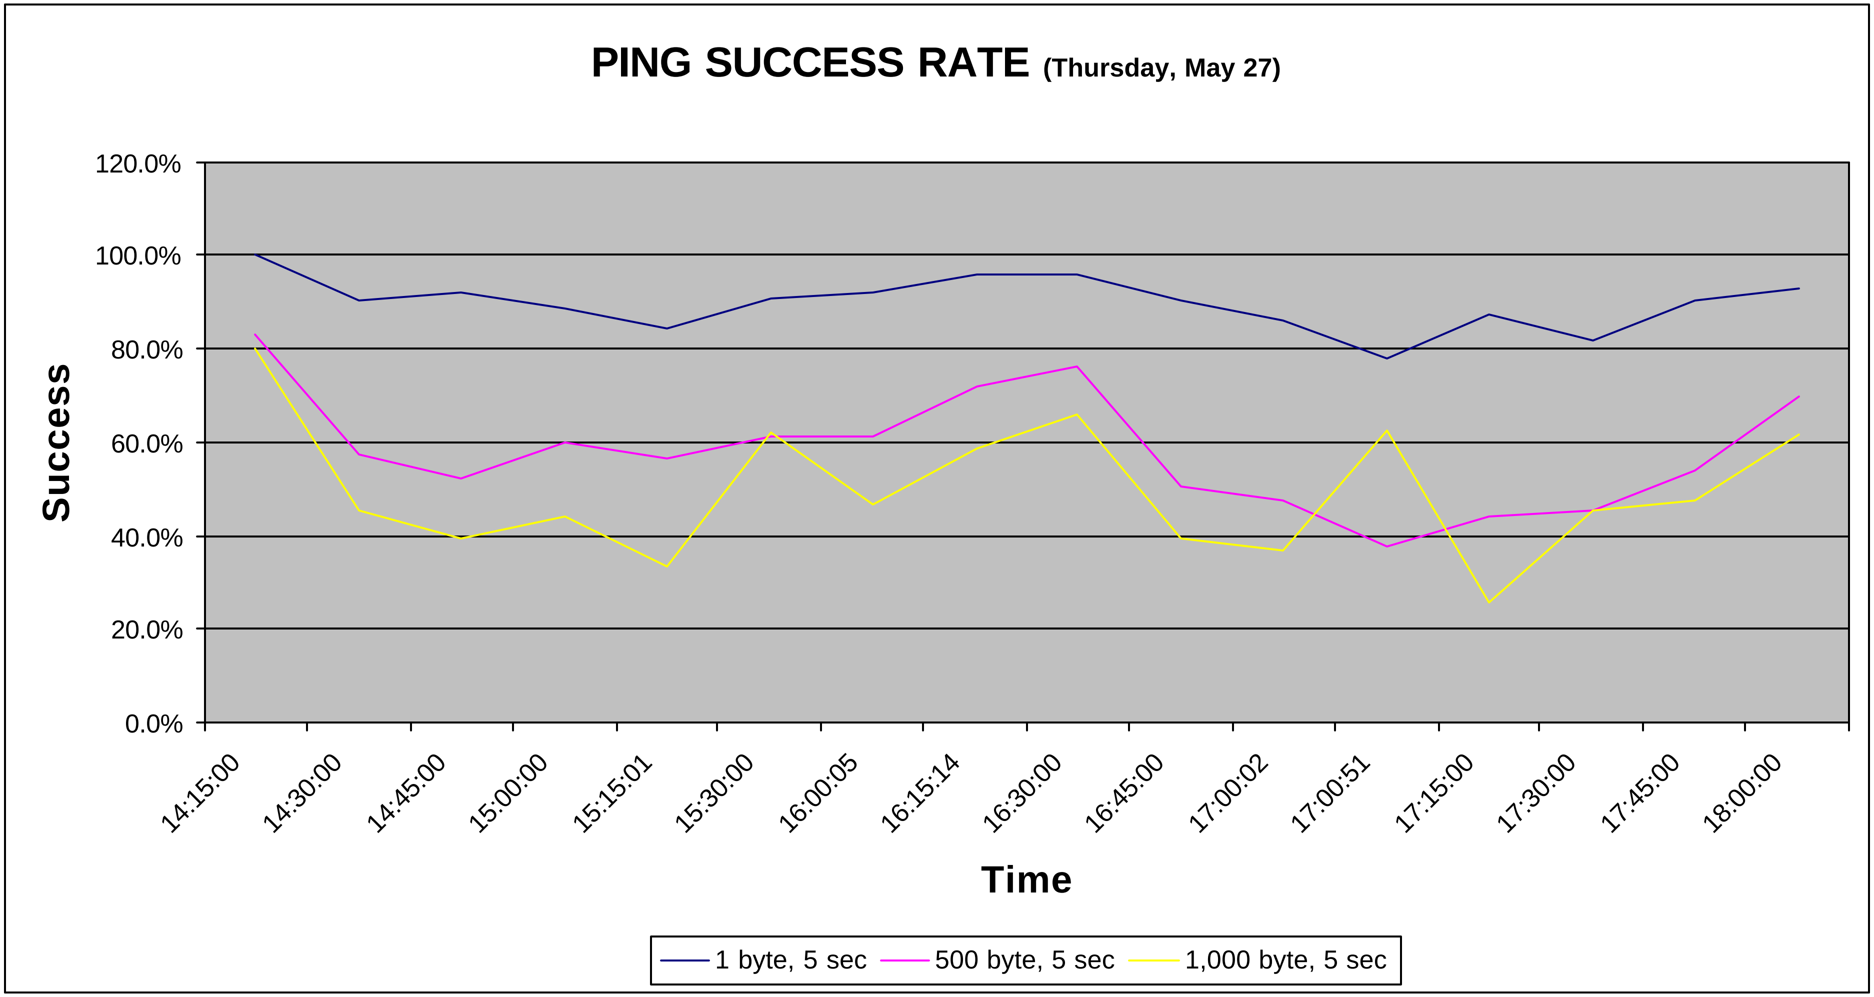
PING SUCCESS RATE (Thursday, May 27)
Success
120.0%
100.0%
80.0%
60.0%
40.0%
20.0%
0.0%
14:15:00
14:30:00
14:45:00
15:00:00
15:15:01
15:30:00
16:00:05
16:15:14
16:30:00
16:45:00
17:00:02
17:00:51
17:15:00
17:30:00
17:45:00
18:00:00
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec

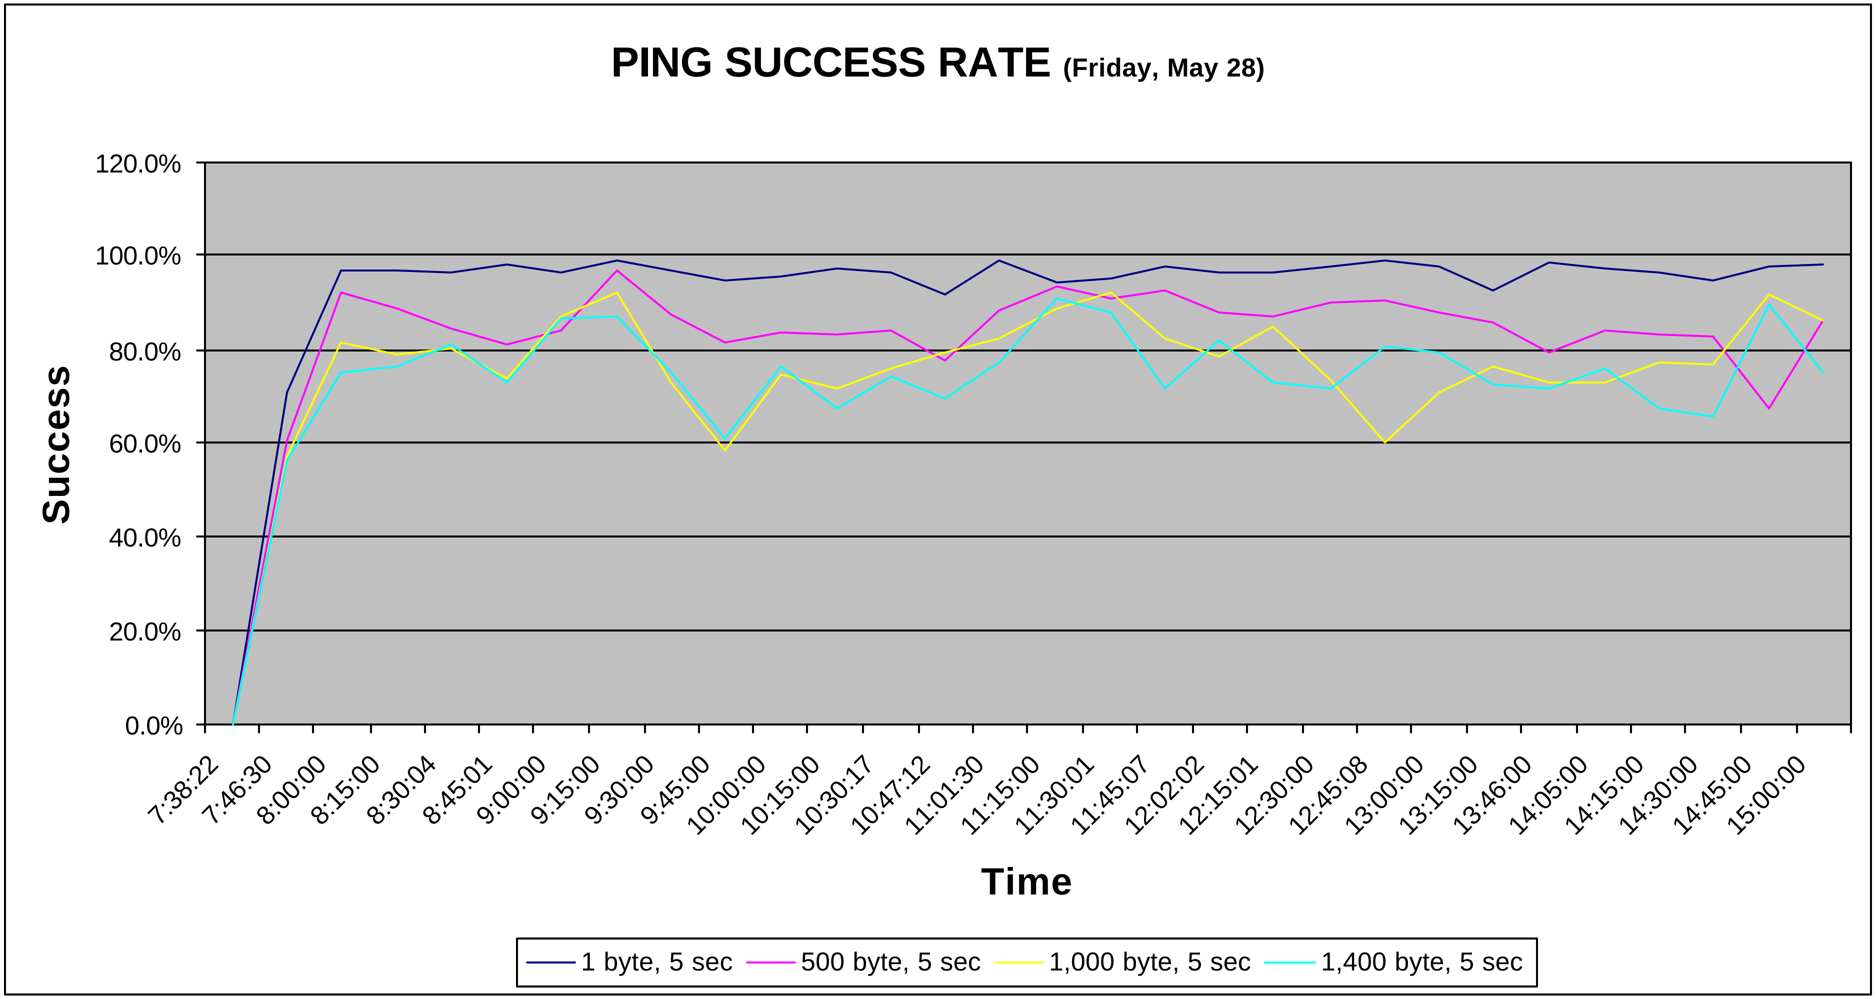
PING SUCCESS RATE (Friday, May 28)
120.0%
100.0%
80.0%
60.0%
40.0%
20.0%
0.0%
Success
7:38:22
7:46:30
8:00:00
8:15:00
8:30:04
8:45:01
9:00:00
9:15:00
9:30:00
9:45:00
10:00:00
10:15:00
10:30:17
10:47:12
11:01:30
11:15:00
11:30:01
11:45:07
12:02:02
12:15:01
12:30:00
12:45:08
13:00:00
13:15:00
13:46:00
14:05:00
14:15:00
14:30:00
14:45:00
15:00:00
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec
1,400 byte, 5 sec

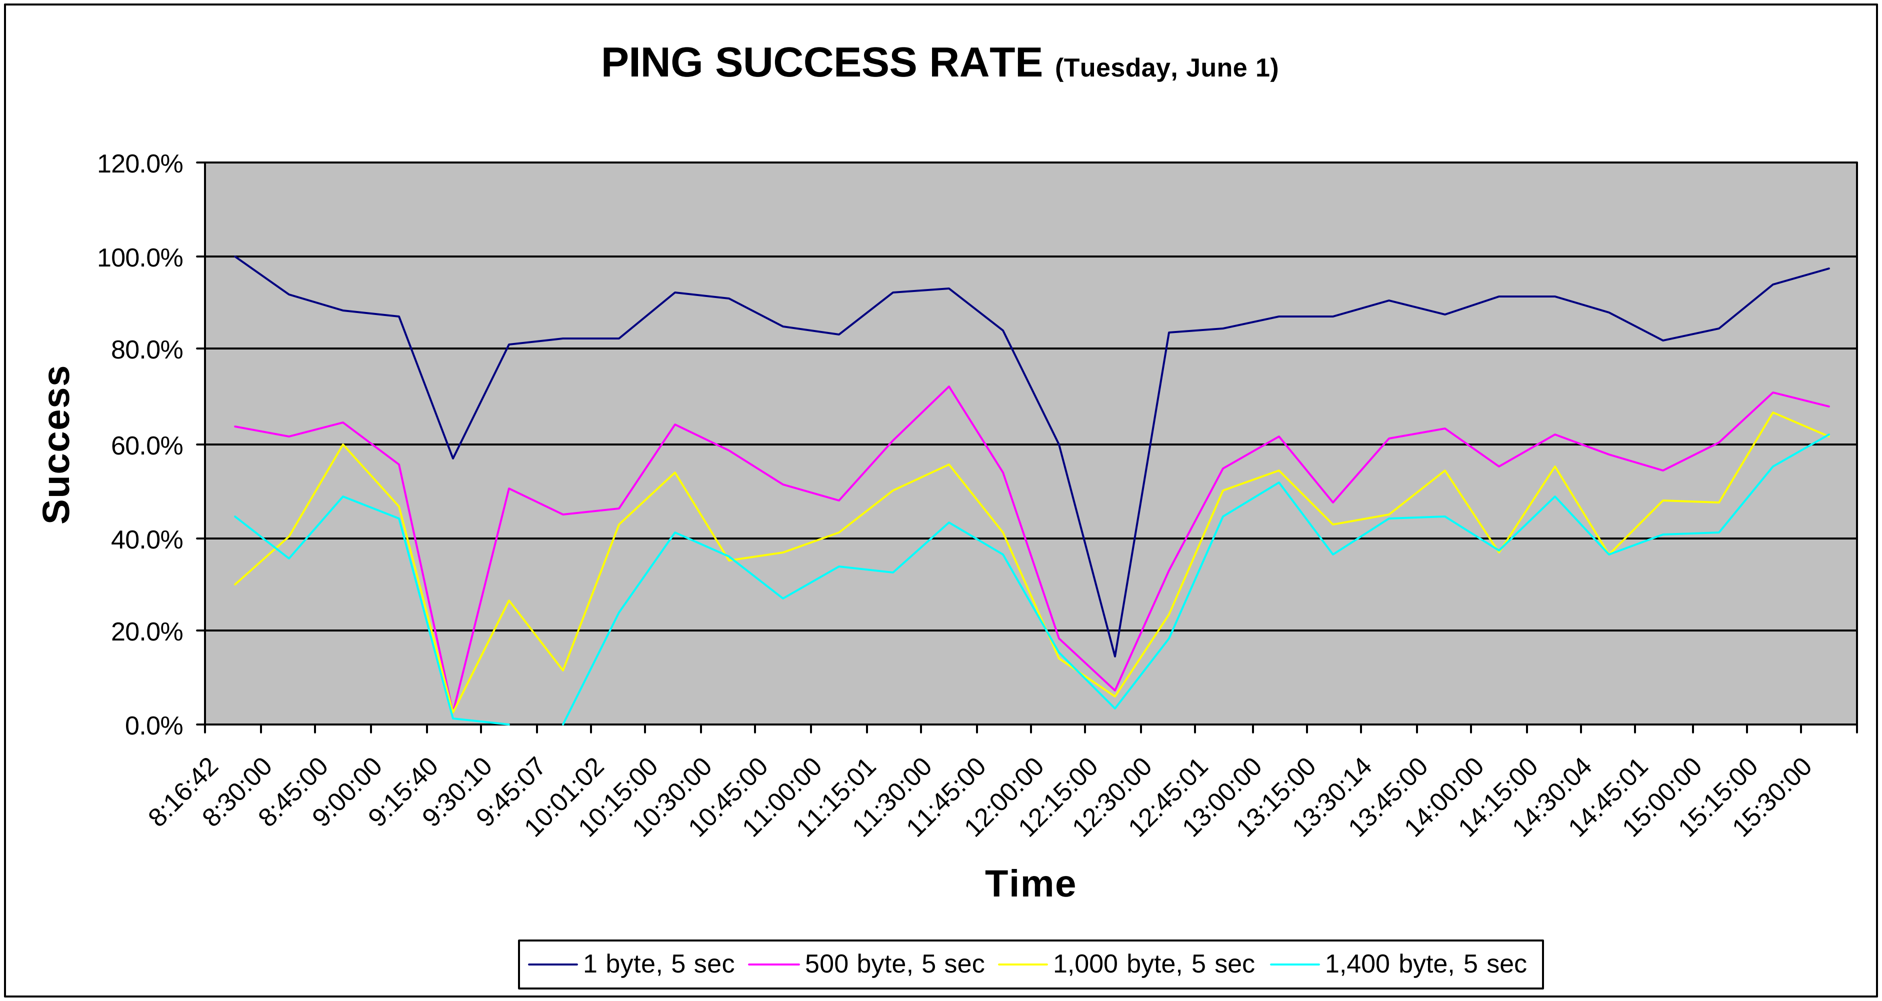
PING SUCCESS RATE (Tuesday, June 1)
120.0%
100.0%
80.0%
60.0%
40.0%
20.0%
0.0%
Success
8:16:42
8:30:00
8:45:00
9:00:00
9:15:40
9:30:10
9:45:07
10:01:02
10:15:00
10:30:00
10:45:00
11:00:00
11:15:01
11:30:00
11:45:00
12:00:00
12:15:00
12:30:00
12:45:01
13:00:00
13:15:00
13:30:14
13:45:00
14:00:00
14:15:00
14:30:04
14:45:01
15:00:00
15:15:00
15:30:00
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec
1,400 byte, 5 sec

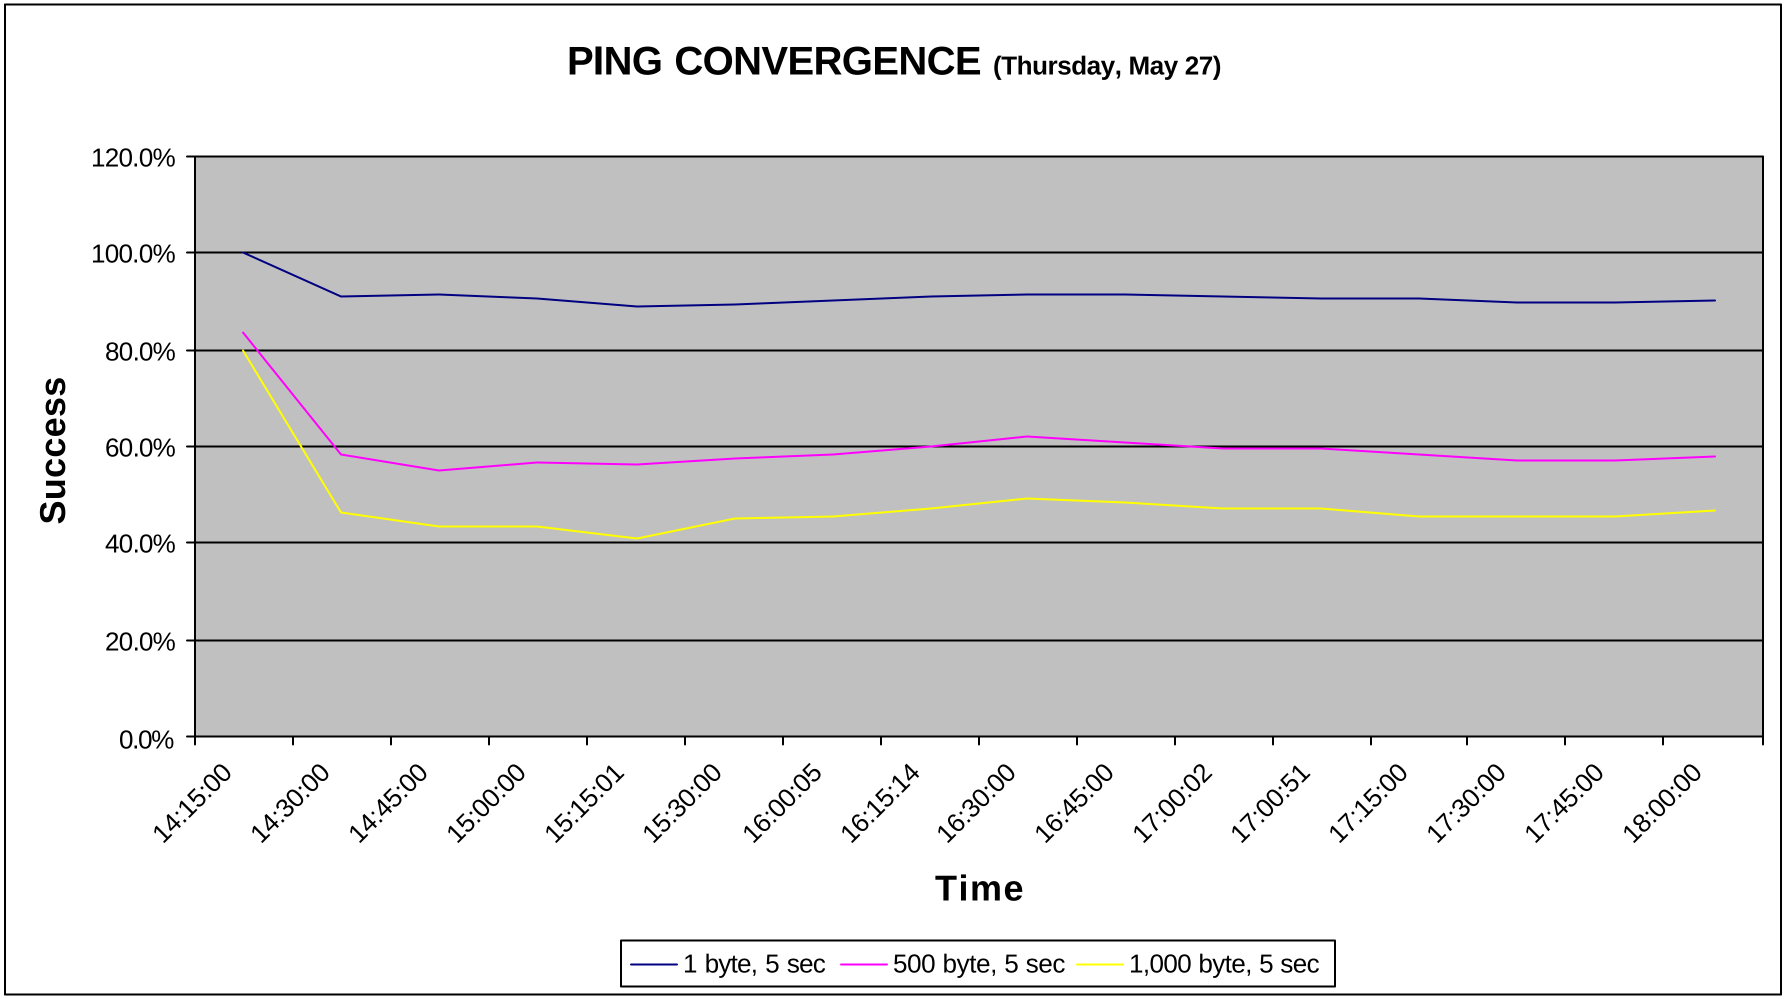
PING CONVERGENCE (Thursday, May 27)
120.0%
100.0%
80.0%
60.0%
40.0%
20.0%
0.0%
Success
14:15:00
14:30:00
14:45:00
15:00:00
15:15:01
15:30:00
16:00:05
16:15:14
16:30:00
16:45:00
17:00:02
17:00:51
17:15:00
17:30:00
17:45:00
18:00:00
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec

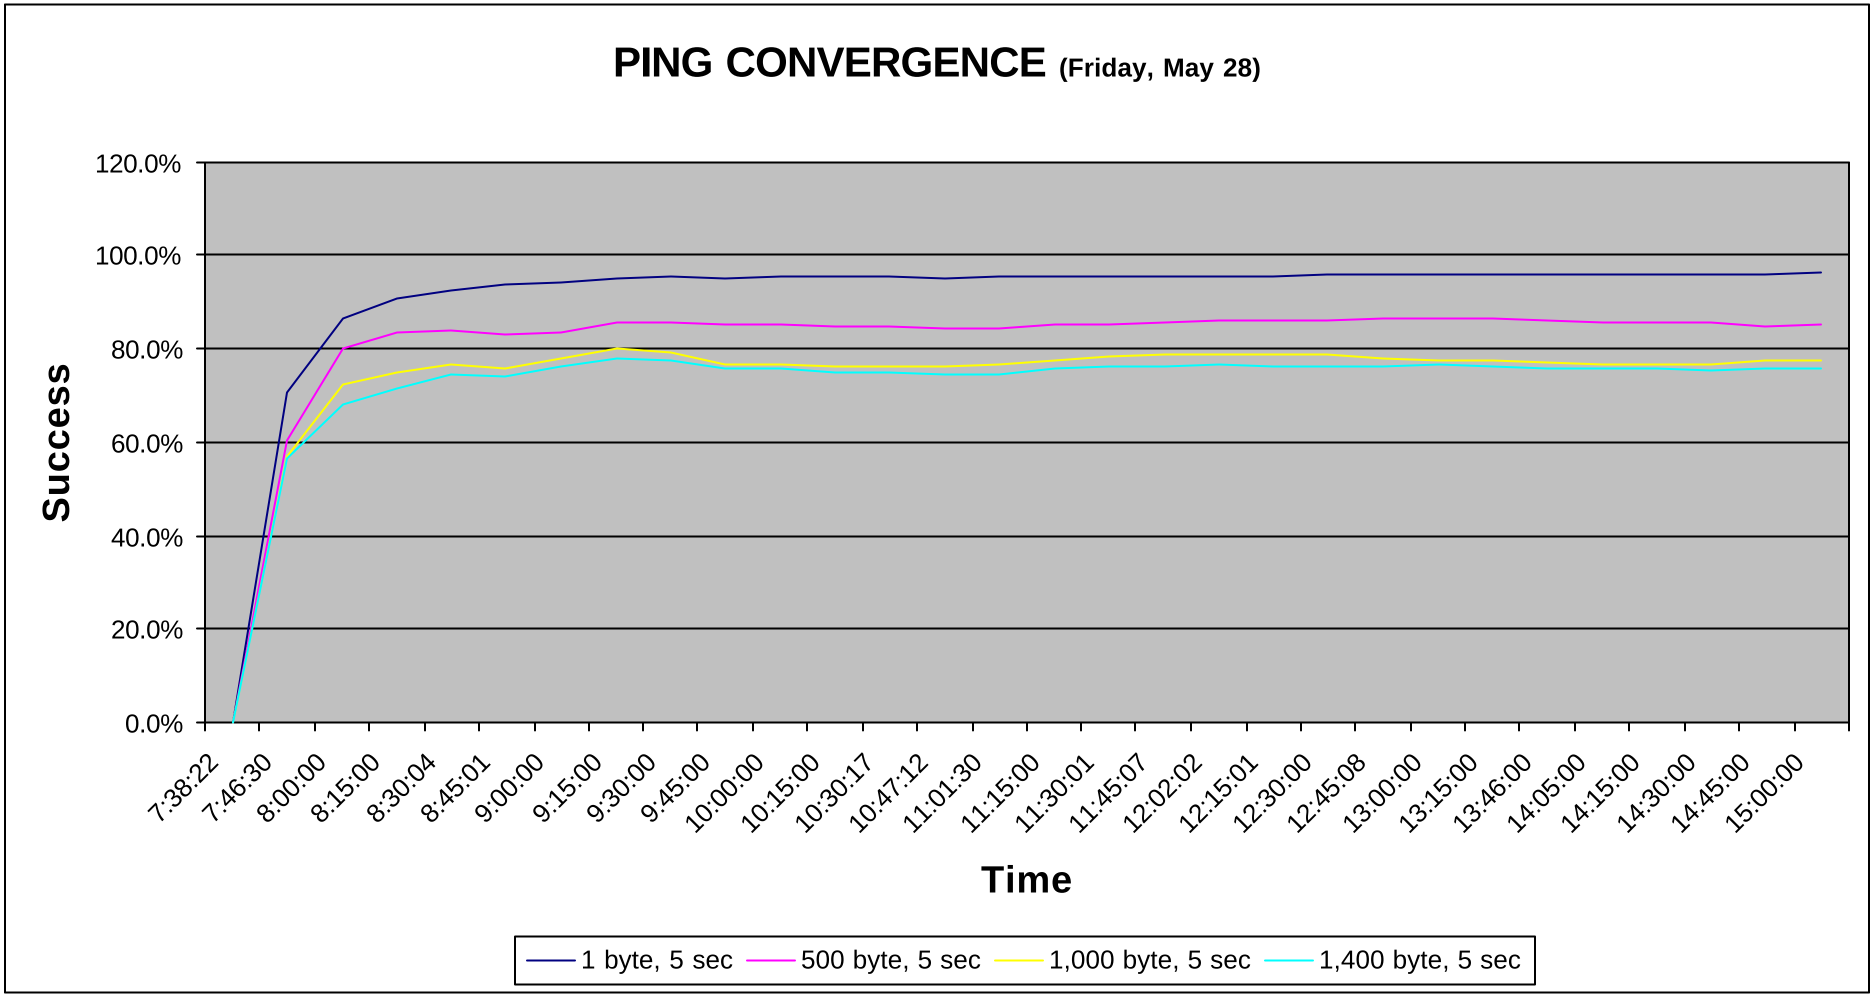
PING CONVERGENCE (Friday, May 28)
Success
120.0%
100.0%
80.0%
60.0%
40.0%
20.0%
0.0%
7:38:22
7:46:30
8:00:00
8:15:00
8:30:04
8:45:01
9:00:00
9:15:00
9:30:00
9:45:00
10:00:00
10:15:00
10:30:17
10:47:12
11:01:30
11:15:00
11:30:01
11:45:07
12:02:02
12:15:01
12:30:00
12:45:08
13:00:00
13:15:00
13:46:00
14:05:00
14:15:00
14:30:00
14:45:00
15:00:00
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec
1,400 byte, 5 sec

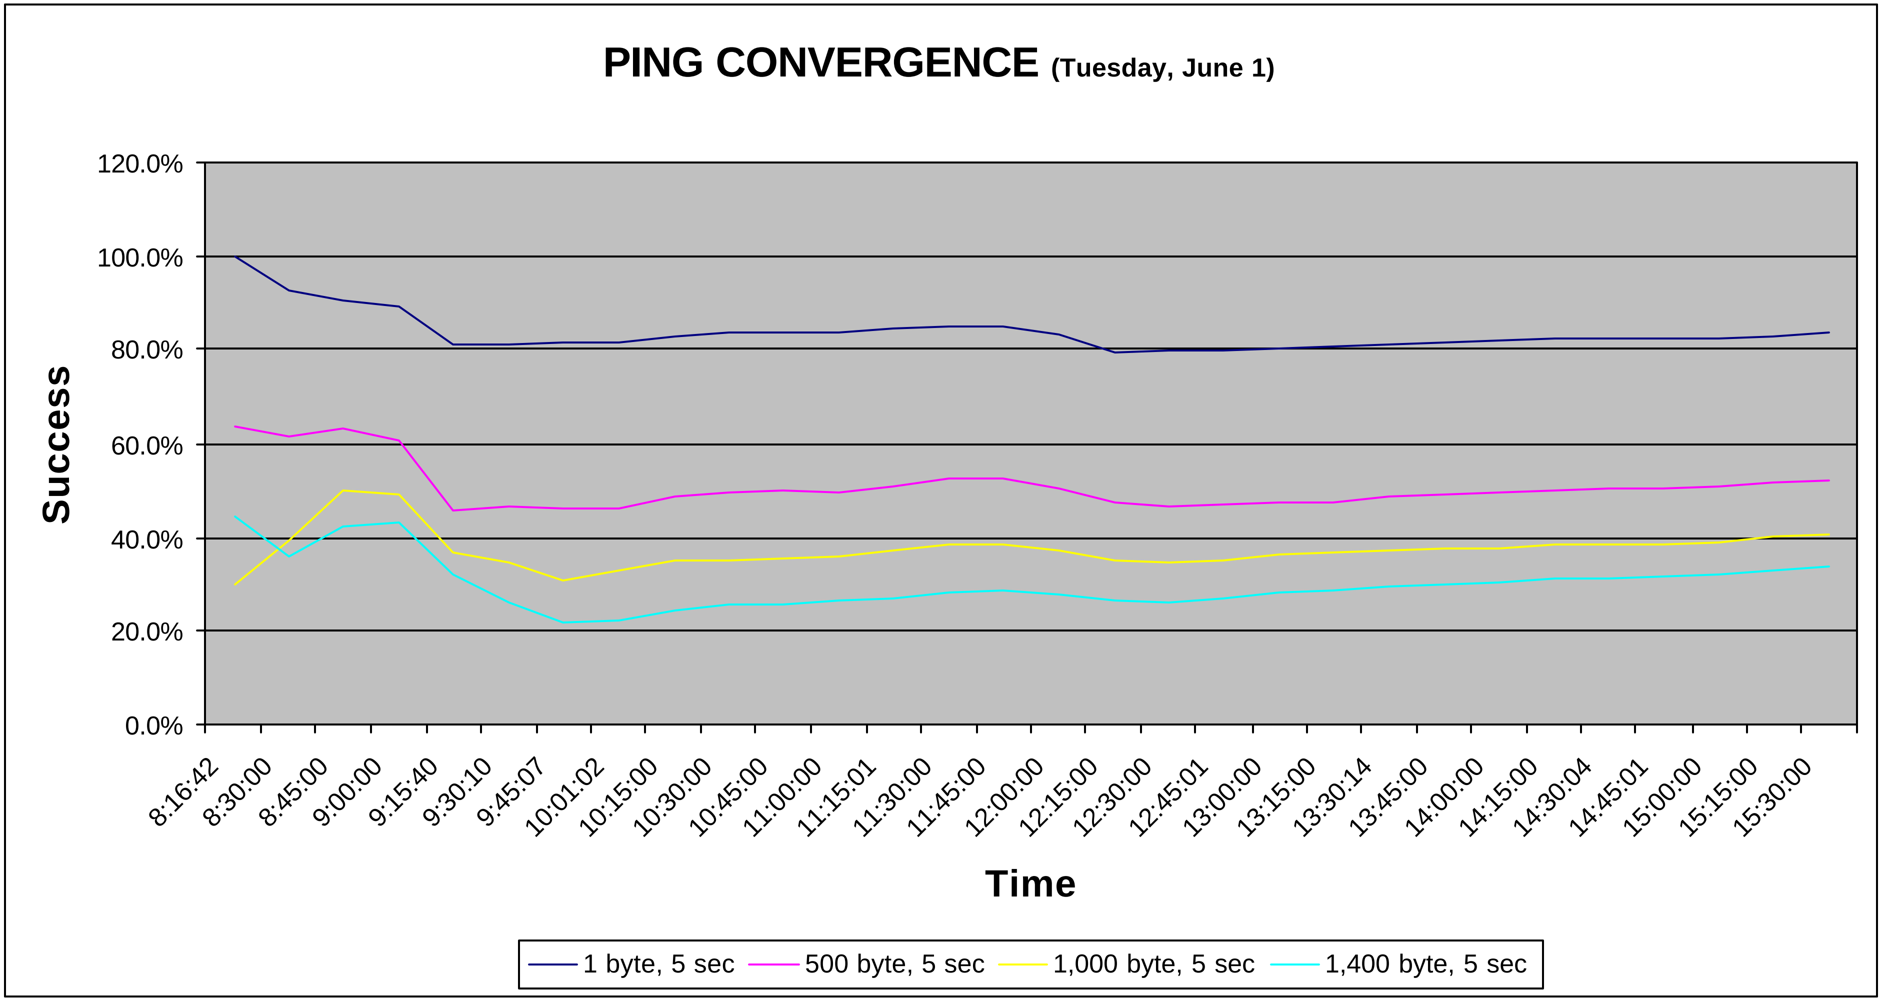
PING CONVERGENCE (Tuesday, June 1)
Success
120.0%
100.0%
80.0%
60.0%
40.0%
20.0%
0.0%
8:16:42
8:30:00
8:45:00
9:00:00
9:15:40
9:30:10
9:45:07
10:01:02
10:15:00
10:30:00
10:45:00
11:00:00
11:15:01
11:30:00
11:45:00
12:00:00
12:15:00
12:30:00
12:45:01
13:00:00
13:15:00
13:30:14
13:45:00
14:00:00
14:15:00
14:30:04
14:45:01
15:00:00
15:15:00
15:30:00
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec
1,400 byte, 5 sec

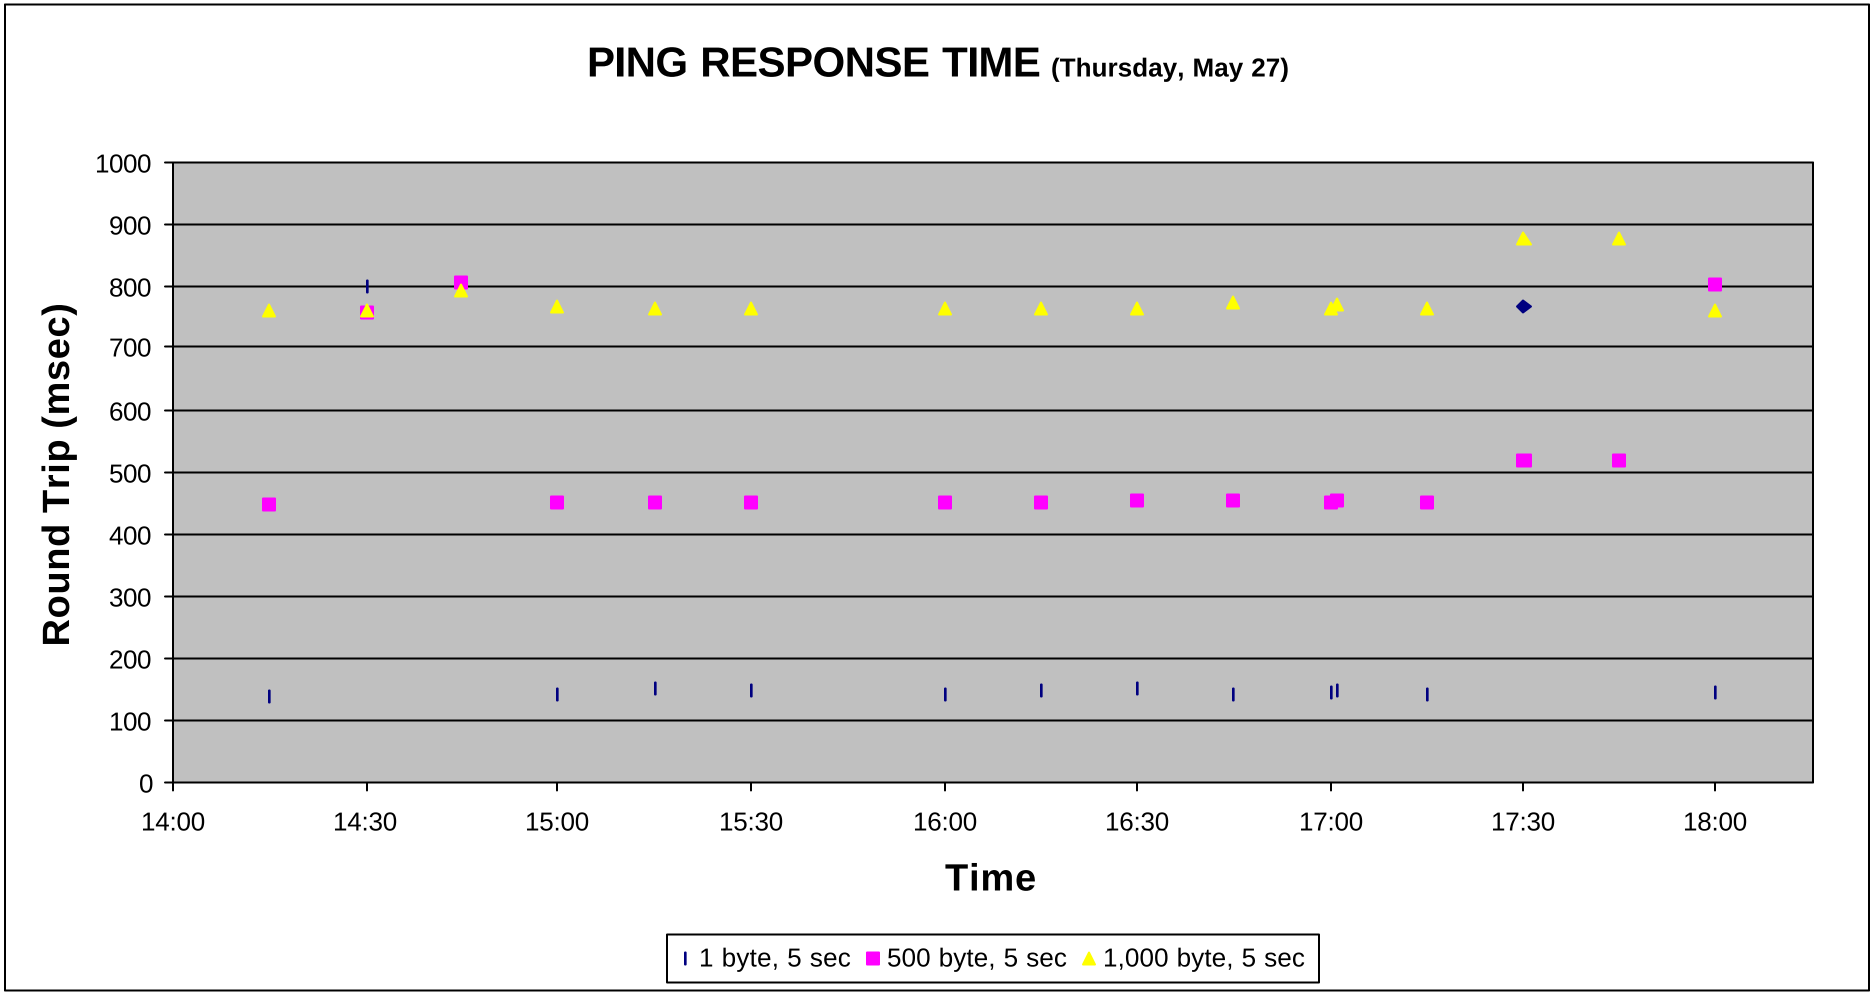
PING RESPONSE TIME (Thursday, May 27)
Round Trip (msec)
0
100
200
300
400
500
600
700
800
900
1000
14:00
14:30
15:00
15:30
16:00
16:30
17:00
17:30
18:00
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec

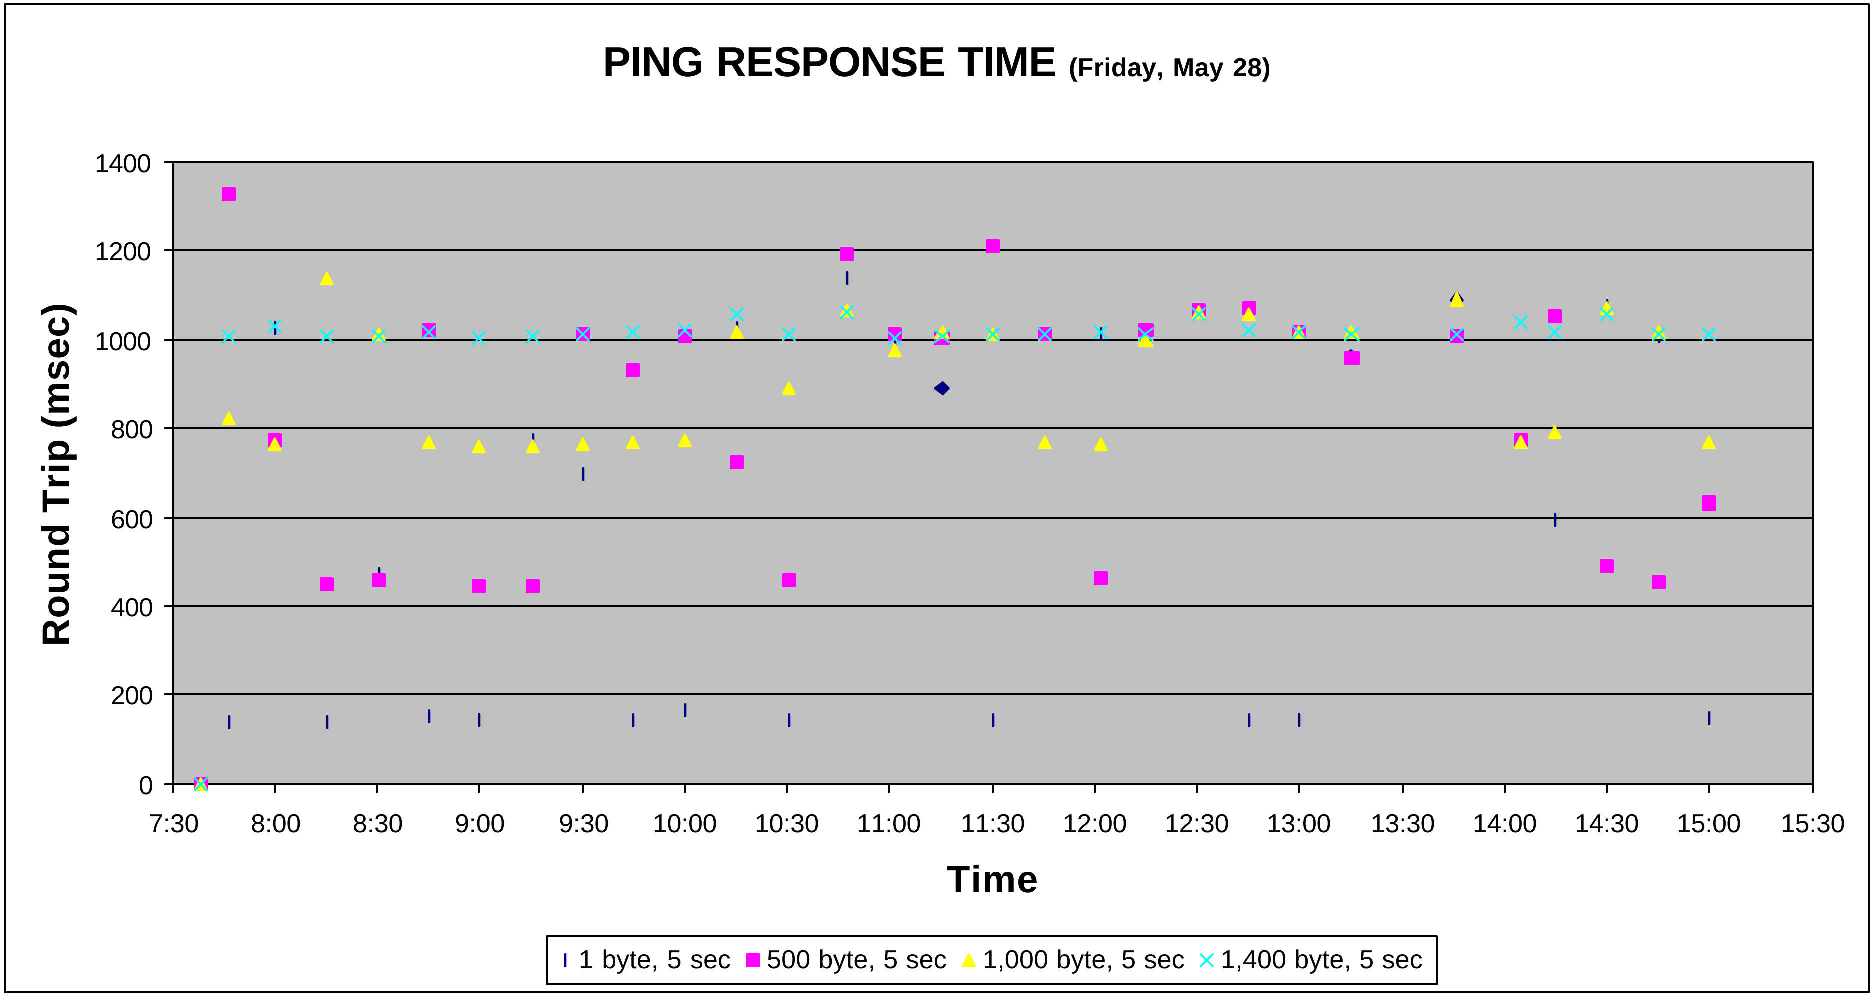

PING RESPONSE TIME (Friday, May 28)
Round Trip (msec)
1400
1200
1000
800
600
400
200
0
7:30
8:00
8:30
9:00
9:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
14:00
14:30
15:00
15:30
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec
1,400 byte, 5 sec

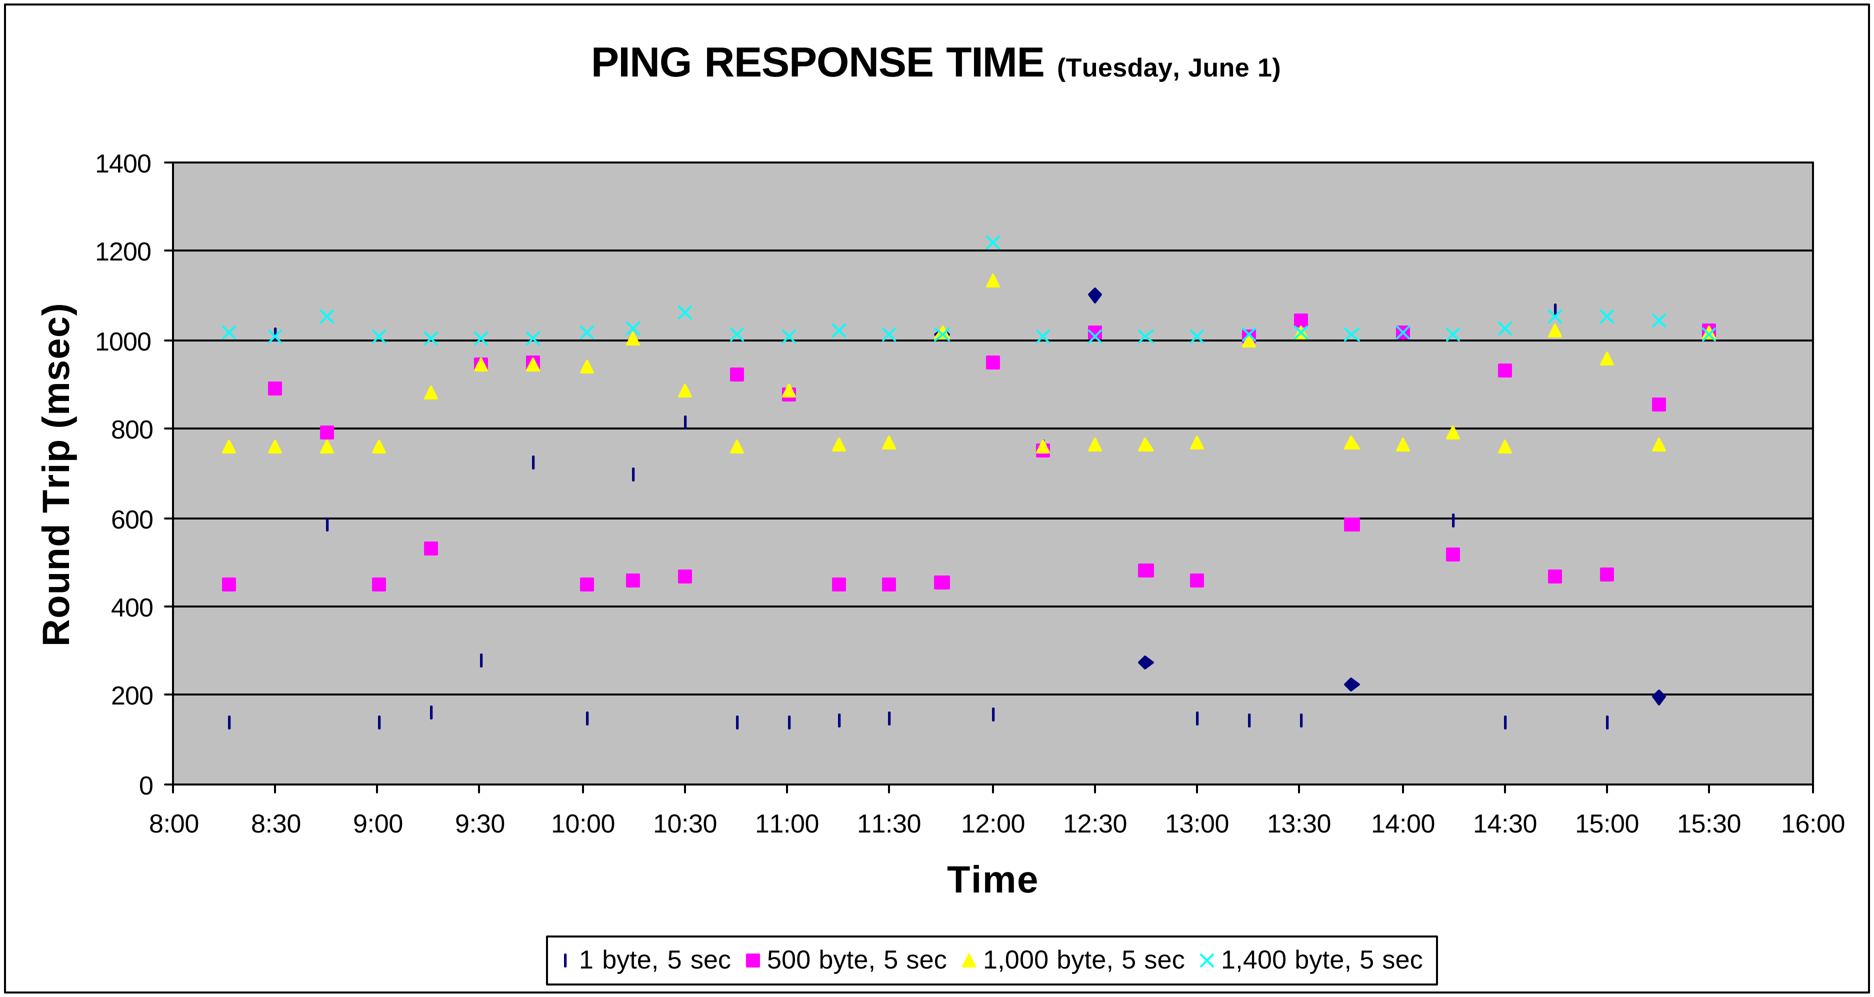
PING RESPONSE TIME (Tuesday, June 1)
Round Trip (msec)
1400
1200
1000
800
600
400
200
0
8:00
8:30
9:00
9:30
10:00
10:30
11:00
11:30
12:00
12:30
13:00
13:30
14:00
14:30
15:00
15:30
16:00
Time
1 byte, 5 sec
500 byte, 5 sec
1,000 byte, 5 sec
1,400 byte, 5 sec

Appendix C

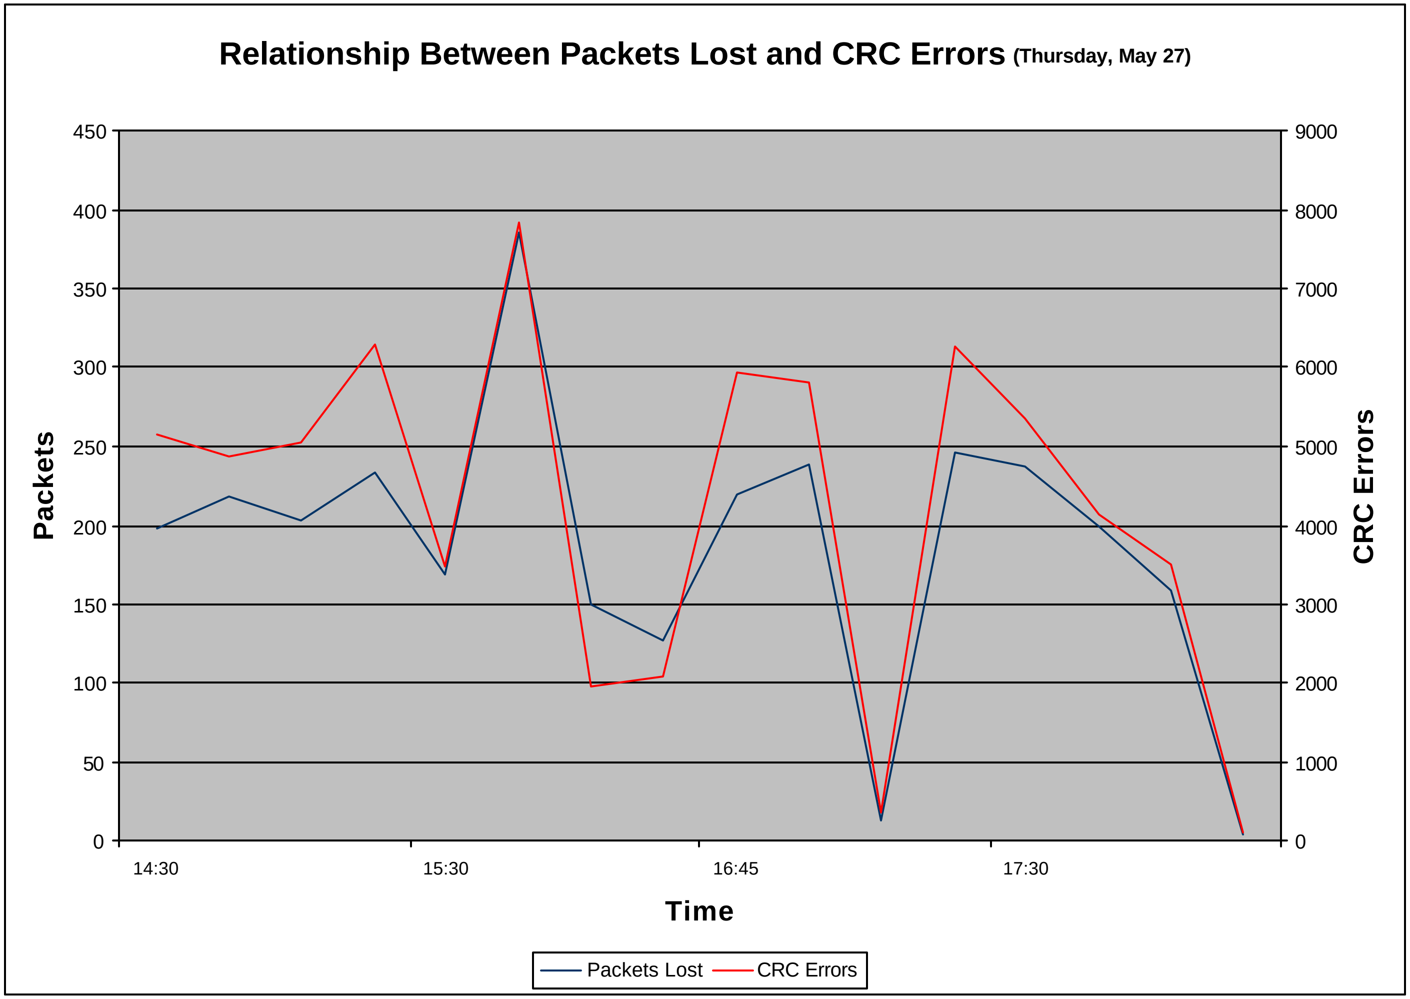

Relationship Between Packets Lost and CRC Errors (Thursday, May 27)
Packets
CRC Errors
0
50
100
150
200
250
300
350
400
450
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
14:30
15:30
16:45
17:30
Time
Packets Lost
CRC Errors

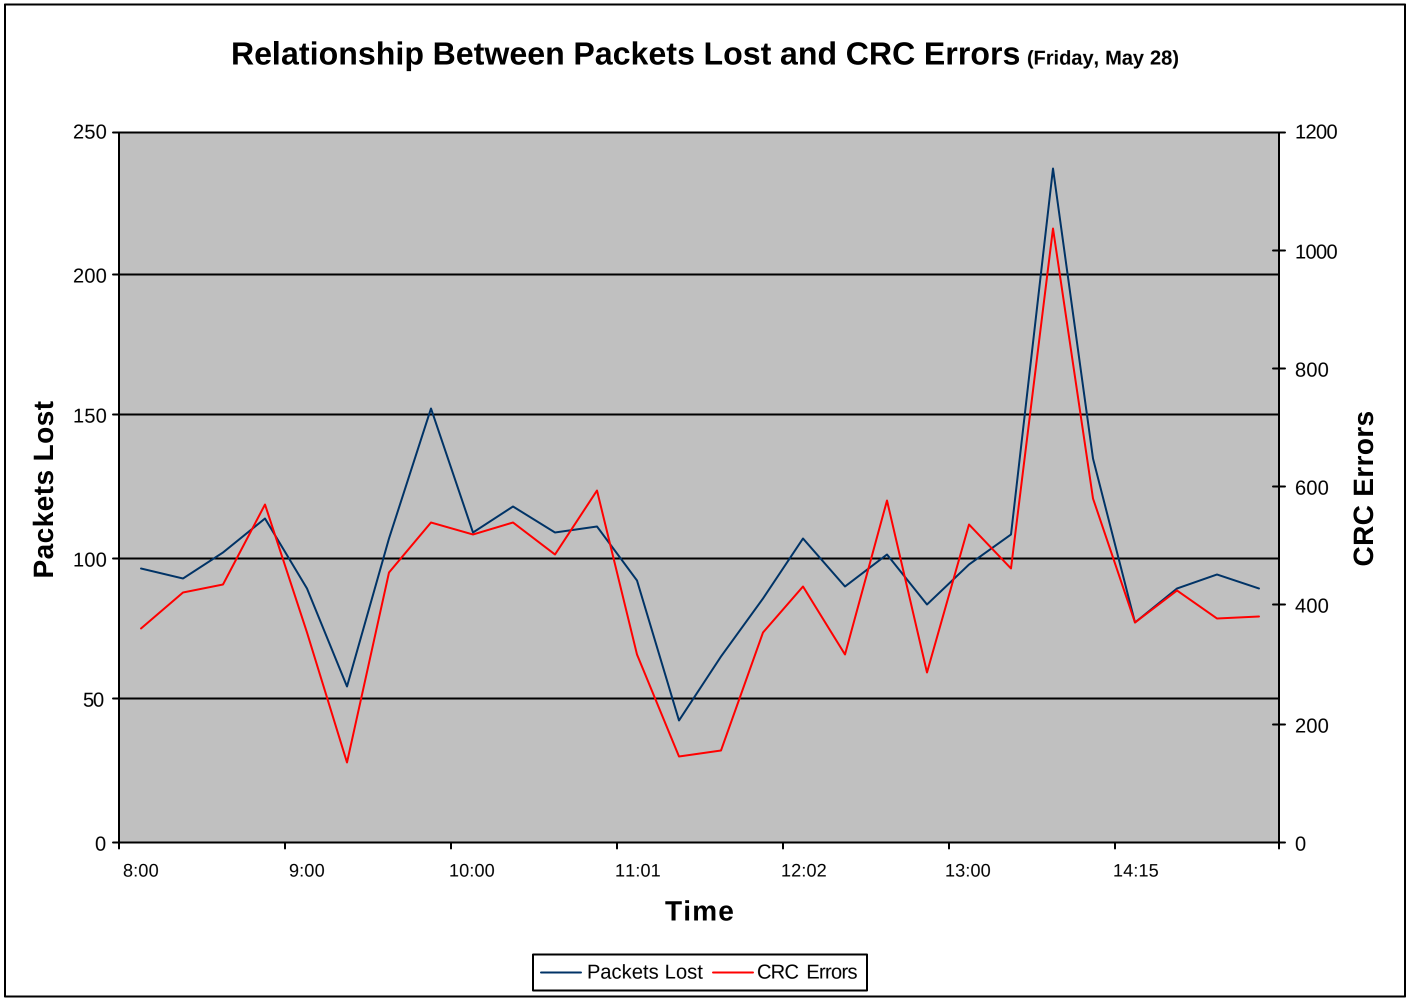
Relationship Between Packets Lost and CRC Errors (Friday, May 28)
Packets Lost
CRC Errors
0
50
100
150
200
250
0
200
400
600
800
1000
1200
8:00
9:00
10:00
11:01
12:02
13:00
14:15
Time
Packets Lost
CRC Errors

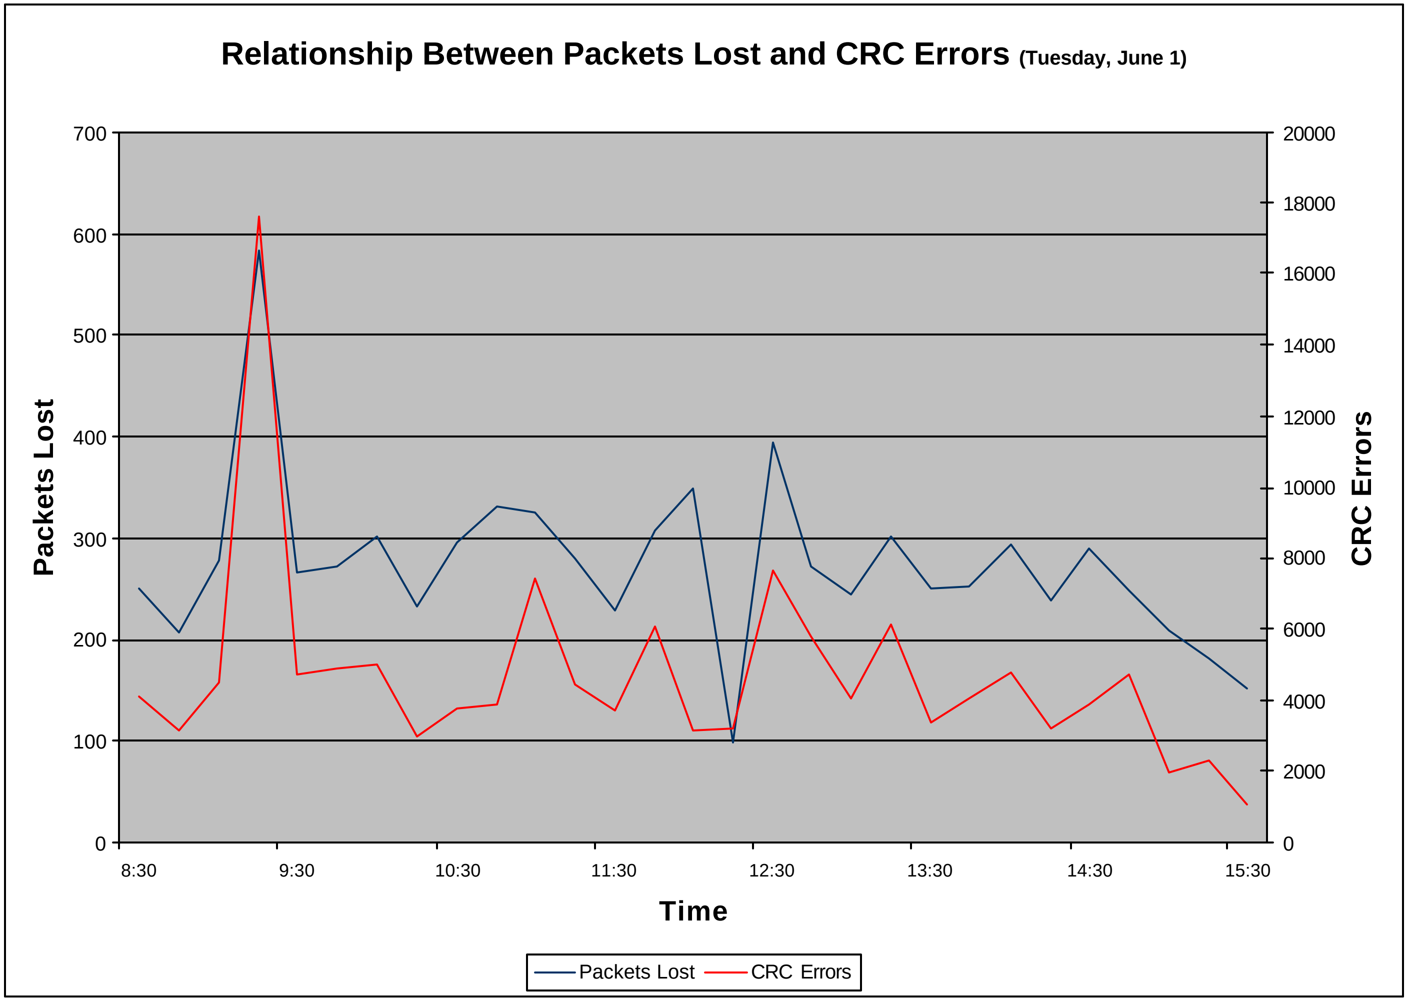
Relationship Between Packets Lost and CRC Errors (Tuesday, June 1)
Packets Lost
CRC Errors
Time
700
600
500
400
300
200
100
0
20000
18000
16000
14000
12000
10000
8000
6000
4000
2000
0
8:30
9:30
10:30
11:30
12:30
13:30
14:30
15:30
Packets Lost
CRC Errors